

WHAT IS CLAIMED IS:

1. A system for servicing a plurality of communication channels, comprising:

a host adapter configured to service communication requests from communication requestors to communication targets, wherein said host adapter provides for up to a maximum number of communication channels to service the communication requests;

a communication fabric configured to provide a limited bandwidth to said host adapter to service said communication requests, wherein said communication fabric couples said communication targets to said host adapter;

a first memory configured to store a lowest level of a hierarchical channel map, wherein said lowest level comprises one bit for each communication channel supported by said host adapter, wherein each bit of said lowest level is set to indicate that the communication channel to which it is mapped has a pending communication request and is cleared if the communication channel to which it is mapped does not have a pending communication request;

a second memory configured to store a top level of the hierarchical channel map, wherein each bit of said top level maps to a section of said lowest level, wherein each bit of said top level is set if at least one bit in the section of said lowest level to which it is mapped is set and is cleared if none of the bits in the section of said lowest level to which it is mapped are set;

wherein said host adapter is configured to determine a next channel to be serviced by examining said hierarchical channel map.

1 2. The system as recited in claim 1, wherein said hierarchical channel map further
2 comprises one or more intermediate levels, wherein each intermediate level comprises a
3 plurality of bits, each one mapping to a group of bits in the next lower intermediate level,
4 wherein each bit indicates if at least one bit in the group to which it maps is set, wherein a
5 highest intermediate level comprises groups of bits mapped to individual one of the bits
6 of said top level, and wherein a lowest intermediate level comprises bits each mapping to
7 one of the groups of said lowest level of the hierarchical channel map.

1 3. The system as recited in claim 2, wherein said host adapter is configured to
2 examine one group of bits at one or more levels of the hierarchical bit map to identify a
3 next channel to be serviced.

1 4. The system as recited in claim 3, further comprising a service mask for each level
2 of the hierarchical channel map, wherein each intermediate level service mask indicates
3 the next bit position to be examined within a selected group of bits to determine a
4 selected group of bits to be examined at the next level, wherein the service mask for the
5 lowest level indicates the next bit position to be examined within a selected group to
6 determine the next channel with a pending request to be services.

1 5. The system as recited in claim 4, wherein said host adapter is configured to
2 service each pending communication request for one service unit before moving on to a
3 next pending communication request, wherein the next pending communication request is
4 determined by examining a current selected group in said hierarchical channel map for
5 the next set bit in order from a position indicated by the corresponding service mask.

1 6. The system as recited in claim 4, wherein said one service unit is less than a
2 maximum message size for the communication requests.

1 7. The system as recited in claim 1, wherein said first memory and said second
2 memory are part of a memory block accessible by said host adapter.

1 8. The system as recited in claim 1, wherein said second memory is a register
2 comprised within an integrated circuit with said host adapter.

1 9. The system as recited in claim 1, further comprising a plurality of service class
2 masks, wherein each service class mask is configured to map a number of said
3 communication channels to one of a plurality of service classes, wherein each service
4 class is allocated a portion of said limited bandwidth on said communication fabric.

1 10. The system as recited in claim 9, further comprising a service array comprising a
2 plurality of entries, wherein each entry indicates one of said service classes to be serviced
3 during a current service unit, wherein said service classes are selected in order according
4 to said entries in said service array, wherein the order is repeated, and wherein the next
5 channel to be serviced is selected from the current service class.

1 11. The system as recited in claim 10, further comprising for each service class a
2 service mask for each level of the hierarchical channel map, wherein when its
3 corresponding service class is being serviced each service mask indicates the next bit
4 position to be examined within a selected group of bits to determine a selected group of
5 bits to be examined at the next level, except wherein the service mask for the lowest level
6 indicates the next bit position to be examined within a selected group to determine the
7 next channel with a pending request to be services.

1 ~~12.~~ A method for servicing a plurality of communication channels, comprising:

2

3 for a first service time:

4

5 selecting a set bit in a top level of a hierarchical channel map, wherein the

6 set bit in the top level indicates a group of bits in a next level of the
7 hierarchical channel map to examine;

8
9 examining in each of one or more intermediate levels of the hierarchical
10 bit map only a group of bits indicated by the set bit selected in the
11 previous level and selecting a set bit from each examined group,
12 wherein said top level is the previous level for a first intermediate
13 level;

14
15 examining in a lowest level of the hierarchical bit map only a group of bits
16 indicated by the previous intermediate level and selecting a set bit
17 from the examined group of the lowest level, wherein the selected
18 bit at the lowest level indicates one of the plurality of
19 communication channels to be serviced; and

20
21 servicing a communication request from the communication channel
22 indicated by the selected bit from the lowest level of the
23 hierarchical channel map.

1 13. The method as recited in claim 12, further comprising:

2
3 for subsequent service times:

4
5 selecting the next set bit from the group of bits in the lowest level from
6 which the previous set bit was selected, or if no more bits are set in
7 that group, selecting the next set bit from the next higher level in
8 the same group as the previous selected bit at that level or
9 repeating said selecting the next set bit from the next higher level
10 until a set bit is found and then selecting a set bit from each group
11 at each lower level as indicated by the set bit at the previous level;

12 and
13
14 servicing subsequent communication requests during each subsequent
15 communication time from the communication channel indicated by
16 the selected bit from the lowest level of the hierarchical channel
17 map for each subsequent service time.

1 14. The method as recited in claim 12, further comprising:

2
3 setting a bit in the lowest level of the hierarchical channel map for each
4 communication channel that has a pending communication request; and
5
6 setting a bit at each higher level of the hierarchical bit map to indicate that a
7 corresponding group of bits in the next lower level has at least one set bit.

1 15. The method as recited in claim 12, further comprising updating a service mask for
2 each level of the hierarchical channel map to indicate the position of the last selected set
3 bit within the indicated group of bits, wherein each time an indicated group of bits is
4 examined to select the next set bit, it is examined starting at the next bit position from the
5 position indicated by the corresponding service mask

1 16. The method as recited in claim 13, wherein each bit position of the top level of the
2 hierarchical channel map represents a different section of the plurality of communication
3 channels, the method further comprising:

4
5 mapping each section of communication channels to a different service class;
6
7 indicating in a service array an order in which each service class is to be serviced;
8
9 choosing a service class according to said service array; and

10

11

examining only portions of the hierarchical channel map corresponding to

12

sections of communication channels mapped to the chosen service class to

13

select the next set bit at the lowest level.

1

17. The method as recited in claim 16, further comprising updating a service mask for

2

each level of the hierarchical channel map and for each service class to indicate the

3

position of the last selected set bit within the indicated group of bits and for the

4

corresponding service class, wherein each time an indicated group of bits is examined to

5

select the next set bit, it is examined starting at the next bit position from the position

6

indicated by the corresponding service mask for the current service class.

1

18. A system for servicing communication queues, comprising:

2

3

memory configured to store a hierarchical channel map comprising a plurality of

4

levels, wherein each bit of the lowest level is mapped to a different one of

5

a plurality of communication channels, and wherein each bit of each

6

higher level is mapped to a group of bits at the next lower level;

7

8

a host adapter configured to maintain the hierarchical bit map, wherein each bit at

9

the lowest level is set if the channel to which it is mapped has a pending

10

communication request and is cleared if not, and wherein each bit of each

11

higher level is set if at least one bit is set in the lower level group to which

12

it is mapped and is cleared if not; and

13

14

wherein said host adapter is configured to examine the hierarchical bit map to

15

determine a next one of the communication channels to service.

1

19. The system as recited in claim 18, wherein said host adapter is configured to

2

service for one service unit a channel request from a channel mapped to a set bit at the

3 lowest level of the hierarchical channel map, wherein the set bit is selected by examining
4 a current group of bits at the lowest level of the hierarchical channel map to select a next
5 set bit in that group indicating a channel with a pending request, and if no more bits are
6 set in the current group, examining a current group at the next higher level to select a next
7 set bit and then examining the next lower level group indicated by the selected higher
8 level set bit.

1 20. The system as recited in claim 19, further comprising a service mask for each
2 level of the hierarchical bit map, wherein each service mask is configured to indicate the
3 next bit position in the current group to be examined for a set bit.

1 21. The system as recited in claim 20, wherein each service mask is configured to
2 indicate the bit position within the current group for the corresponding level of the last
3 selected set bit in that group, wherein the host adapter is configured to examine each
4 group for the next set bit after the bit position indicated by the corresponding service
5 mask.

1 22. The system as recited in claim 18, wherein each group of bits at one level of the
2 hierarchical channel map has the same number of bits.

1 23. The system as recited in claim 18, wherein each group of bits at the lowest level is
2 accessible by a single memory access.

1 24. The system as recited in claim 18, wherein the top level of the hierarchical
2 channel map is stored within a register comprised within the same integrated circuit as
3 said host adapter.

1 25. The system as recited in claim 18, wherein said host adapter is configured to
2 service channel requests according to service classes, wherein each channel is mapped to
3 one or more service classes, and wherein the next channel to be serviced is selected by

4 examining only portions of said hierarchical channel map corresponding to channels
5 mapped to a current service class.

1 26. The system as recited in claim 25, wherein the current service class is selected
2 according to an order of service classes indicated by a service array.

1 27. The system as recited in claim 25, further comprising one service mask per level
2 of the hierarchical bit map per service class, wherein each service mask is configured to
3 indicate the next bit position in the current group to be examined for a set bit when the
4 corresponding service class is the current service class.

1 28. The system as recited in claim 19, wherein said one service unit is a quantum
2 smaller than a maximum message size for the channel requests.

1 ~~29.~~ A computer readable medium comprising program instructions, wherein said
2 program instructions are operable to:

3
4 select a first bit set in a current top level group of a hierarchical channel map after
5 a position indicated by a top level service mask;

6
7 set the top level service mask to indicate the position of the selected first set bit in
8 the current top level group;

9
10 in a second level of the hierarchical channel map, access a second level group
11 indicated by the selected set bit from the top level;

12
13 select a first bit set in the accessed second level group after a position indicated by
14 a second level service mask;

15
16 set the second level service mask to indicate the position of the selected bit in the

17 accessed second level group;
18
19 in a bottom level of the hierarchical channel map, access a bottom level group
20 indicated by the selected set bit from the second level;
21
22 select a first bit set in the accessed bottom level group after a position indicated by
23 a bottom level service mask;
24
25 set the bottom level service mask to indicate the position of the selected bit in the
26 accessed bottom level group; and
27
28 service a request from a channel indicated by the selected set bit from the bottom
29 level.

1 30. A method for servicing a plurality of communication channels, comprising:
2
3 for a first service time, examining a top level of a hierarchical channel map to
4 select a section of the communication channels in which at least one
5 channel has a pending communication request, wherein said top level
6 indicates for each of a plurality of sections of the communication channels
7 if at least channel of that section has a pending communication request;
8
9 for said first service time, examining a portion of one or more intermediate levels
10 of the hierarchical channel map to select a lowest level group of the
11 communication channels in which at least one channel has a pending
12 communication request, wherein each intermediate level indicates for each
13 of a plurality of groups of the communication channels if at least one
14 channel of that group has a pending communication request, wherein the
15 groups at each intermediate level are sized smaller than at the previous
16 intermediate level, and wherein examining a portion of each intermediate

17 level determines which portion of the next hierarchical channel map to
18 examine; and
19
20 selecting a next communication channel to be serviced from the lowest level
21 group selected by said examining a portion of one or more intermediate
22 levels.

1 31. A system for servicing a plurality of communication channels, comprising:
2
3 a first memory configured to store one or more levels of a hierarchical channel
4 map, wherein said one or more levels comprises a lowest level for storing
5 an indication of which ones of the plurality of communication channels
6 have pending communication requests;
7
8 a second memory configured to store a top level of said hierarchical channel map,
9 wherein the plurality of communication channels are organized in channel
10 sections, and wherein for each channel section said top level indicates if at
11 least one of the communication channels within that section has a pending
12 communication request; and
13
14 a host adapter configured to determine a next channel to service by examining
15 said lowest level in said first memory, wherein said host adapter
16 determines the next channel to service by examining no more of said
17 lowest level than a portion of said lowest level corresponding to one of
18 said channel sections indicated by said top level as having at least one
19 pending communication request.

1 32. The system as recited in claim 31, wherein said one or more levels stored in said
2 first memory further comprises an intermediate level of said hierarchical channel map,
3 wherein said communication channels are further organized into channel groups, wherein

4 for each channel group said intermediate level indicates if at least one of the
5 communication channels within that group has a pending communication request, and
6 wherein said host adapter is further configured to determine the next channel to service
7 by examining no more of said lowest level than a portion of said lowest level
8 corresponding to one of said channel groups indicated by said intermediate level as
9 having at least one pending communication request, wherein each channel section
10 includes a plurality of channel groups.

1 33. The system as recited in claim 32, wherein said top level stored in said second
2 memory is further configured to indicate which portions of said intermediate level contain
3 at least one indication that at least one channel within one of said channel groups has a
4 pending communication request.

1 34. The system as recited in claim 31, wherein said lowest level comprises a plurality
2 of low level bits, wherein each low level bit is set if a corresponding one of said plurality
3 of channels has a pending communication request and is cleared if the corresponding one
4 of said plurality of channels does not have a pending communication request.

1 35. The system as recited in claim 34, wherein said one or more levels stored in said
2 first memory further comprises an intermediate level comprising a plurality of
3 intermediate level bits, wherein each intermediate level bit corresponds to a group of low
4 level bits, and wherein each intermediate level bit is set if its corresponding group of low
5 level bits has at least one bit set and is cleared if no bits are set in its corresponding group
6 of low level bits.

1 36. The system as recited in claim 35, wherein said top level comprises a plurality of
2 top level bits, wherein each top level bit corresponds to a group of intermediate level bits,
3 and wherein each top level bit is set if its corresponding group of intermediate level bits
4 has at least one bit set and is cleared if no bits are set in its corresponding group of
5 intermediate level bits.

1 38. The system as recited in claim 31, wherein said second memory is a register
2 comprised within an integrated circuit with said host adapter.

1 38. The system as recited in claim 31, wherein said second memory is a register
2 comprised within an integrated circuit with said host adapter.